Artificial Intelligence
Part I: Machine learning

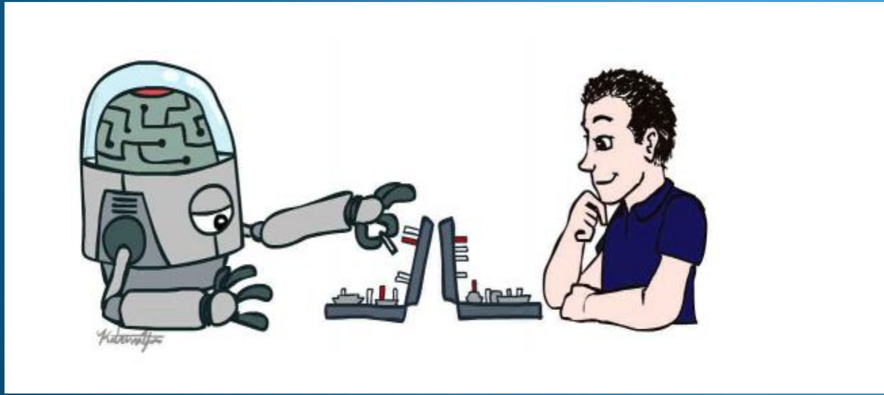

Dr Sherin ElGokhy

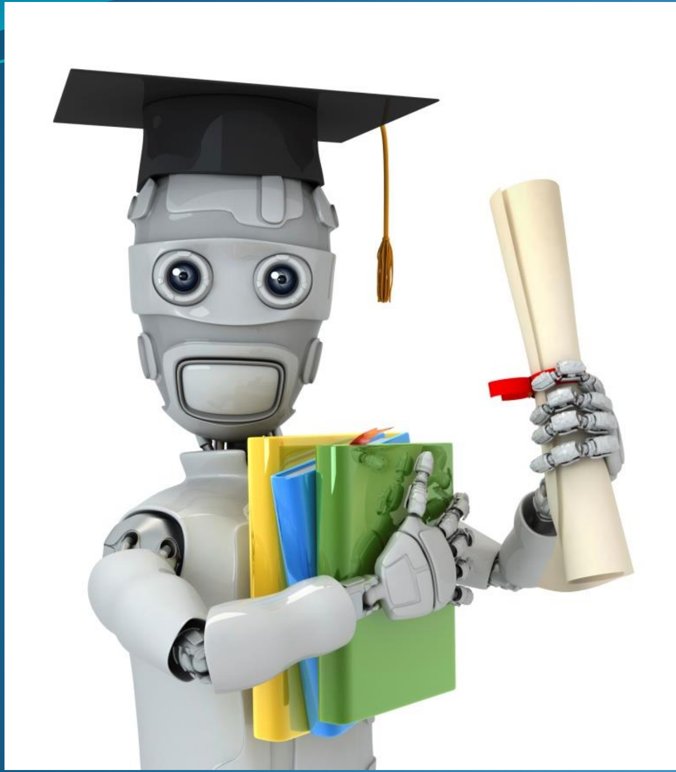

Machine Learning

Clustering

Unsupervised learning introduction

Supervised learning

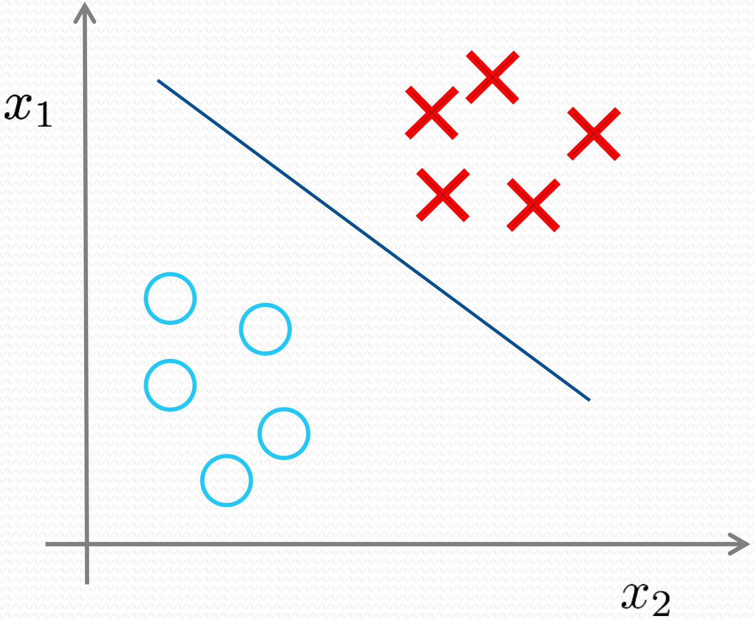

Given labeled training set:

$$\{(x^{(1)}, y^{(1)}), (x^{(2)}, y^{(2)}), (x^{(3)}, y^{(3)}), \ldots, (x^{(m)}, y^{(m)})\}$$

Unsupervised learning

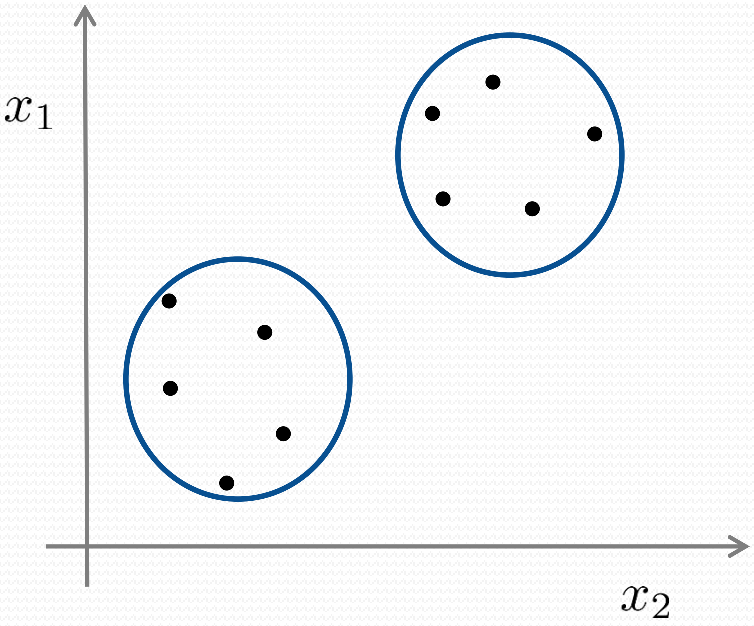

Clustering algorithm: group these data into two different clusters.

Given unlabeled training set: $\{x^{(1)}, x^{(2)}, x^{(3)}, \ldots, x^{(m)}\}$

Applications of clustering

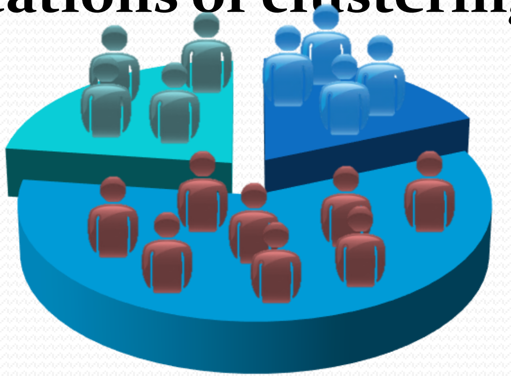

Market segmentation

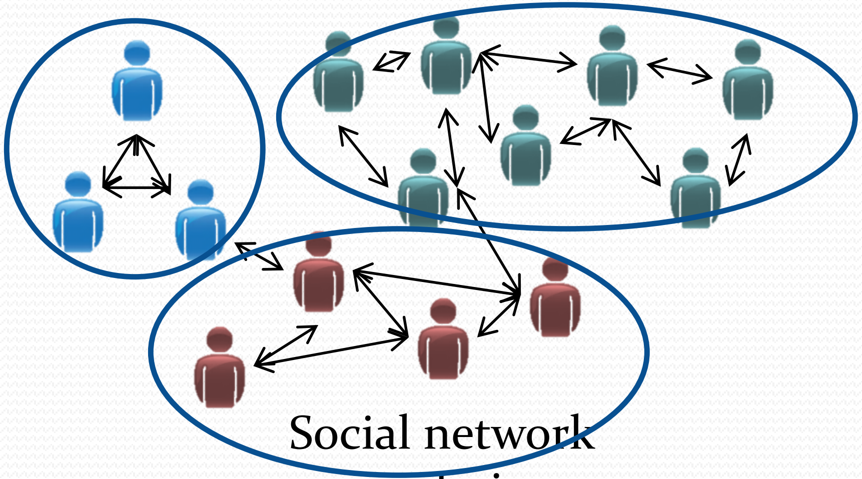

Social network

Organize computing clusters

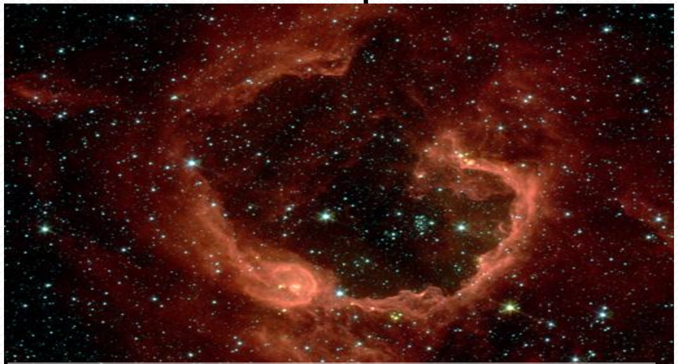

Image credit: NASA/JPL-Caltech/E. Churchwell (Univ. of Wisconsin, Madison)

Astronomical data analysis

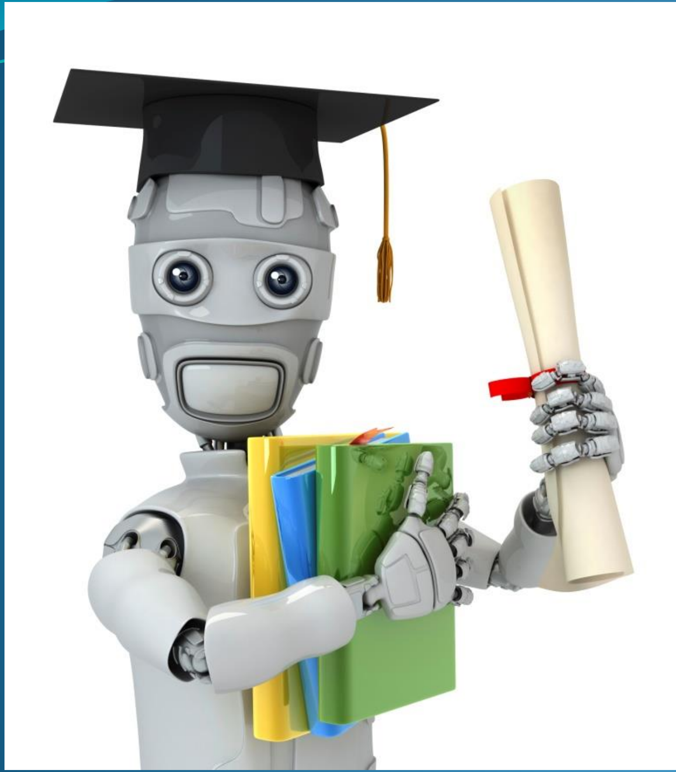

Machine Learning

Clustering

K-means algorithm

- The K Means algorithm is by far the most widely used clustering algorithm.

- K Means is an iterative algorithm and it does two things. First is a cluster assignment step, and second is a move centroid step.

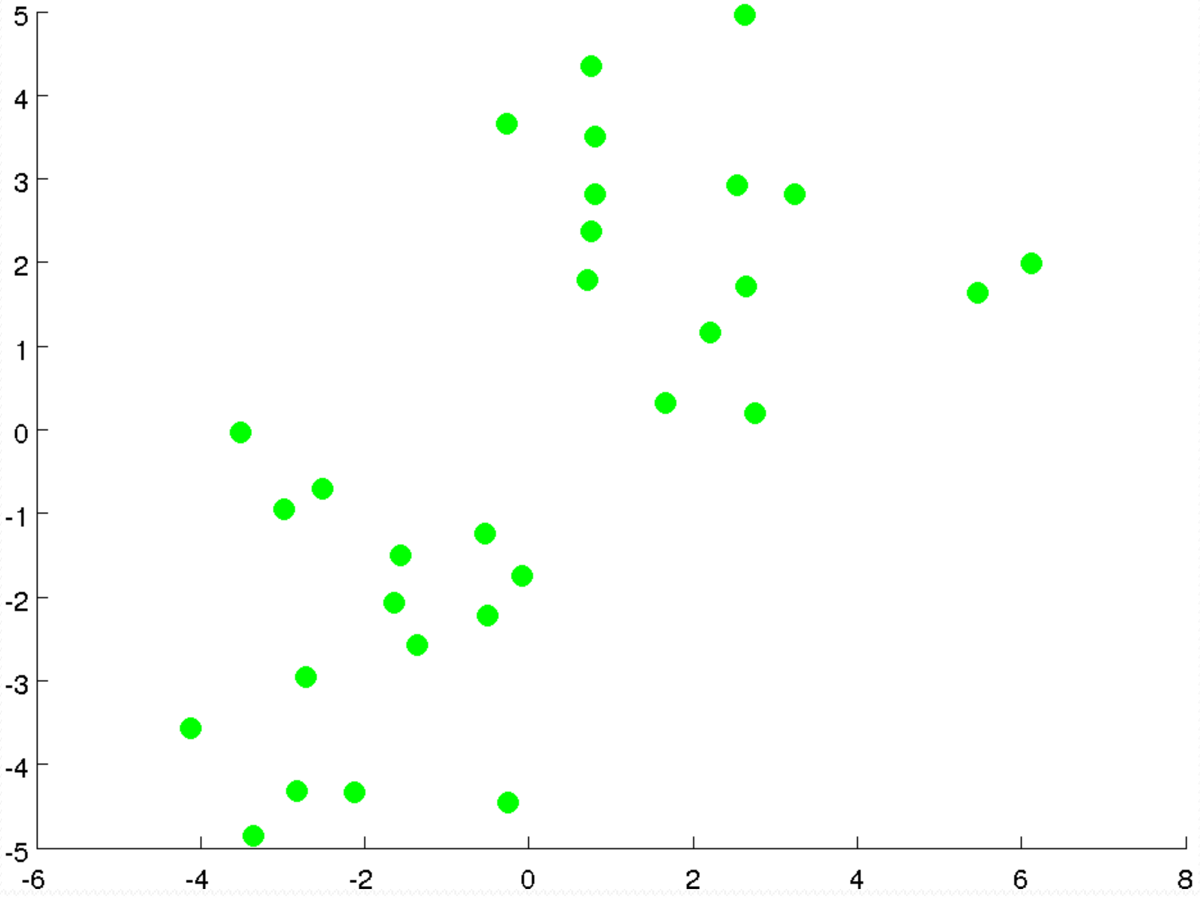
5
4
3
2
1
0
-1
-2
-3
-4
-5
-6
-4
-2
0
2
4
6
8

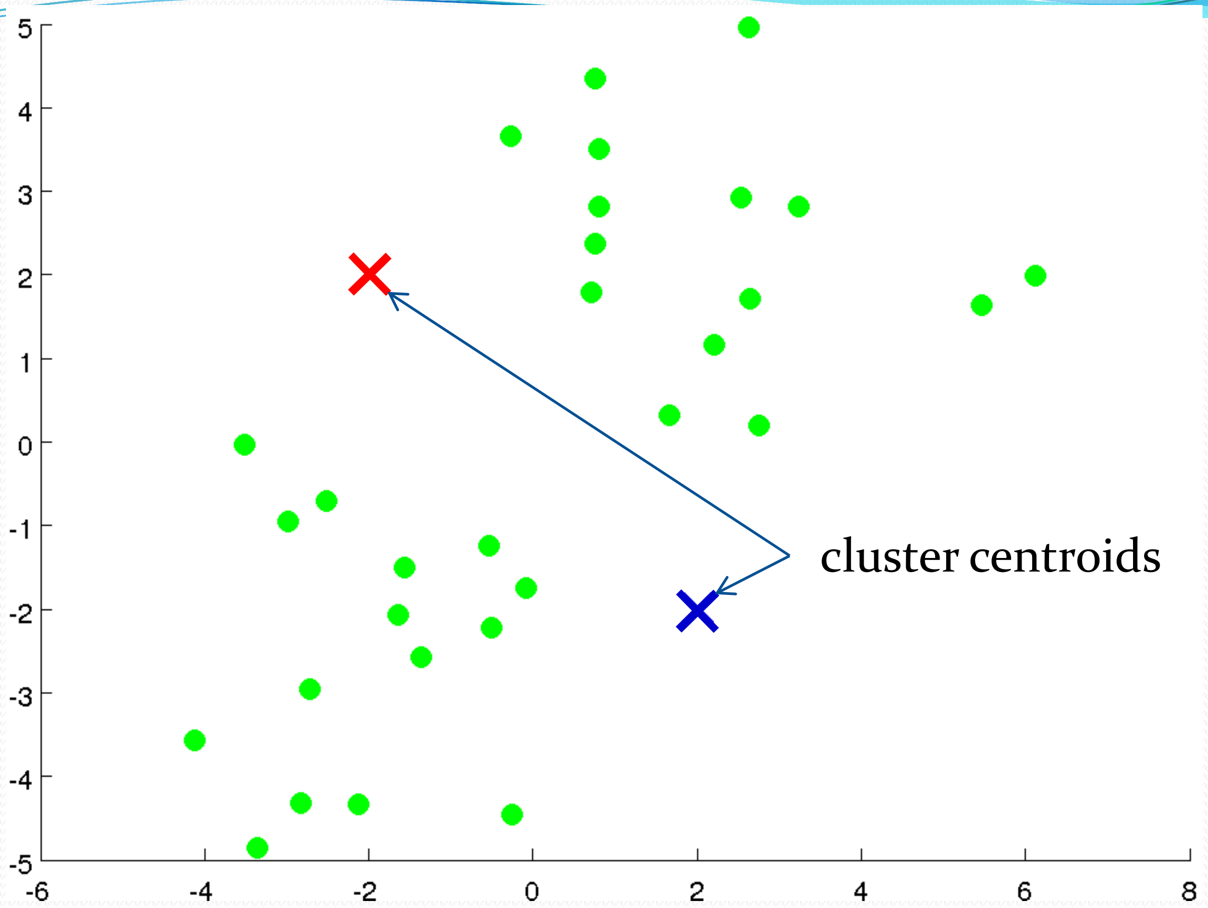

5
4
3
2
1
0
-1
-2
-3
-4
-5
-6
-4
-2
0
2
4
6
8
cluster centroids

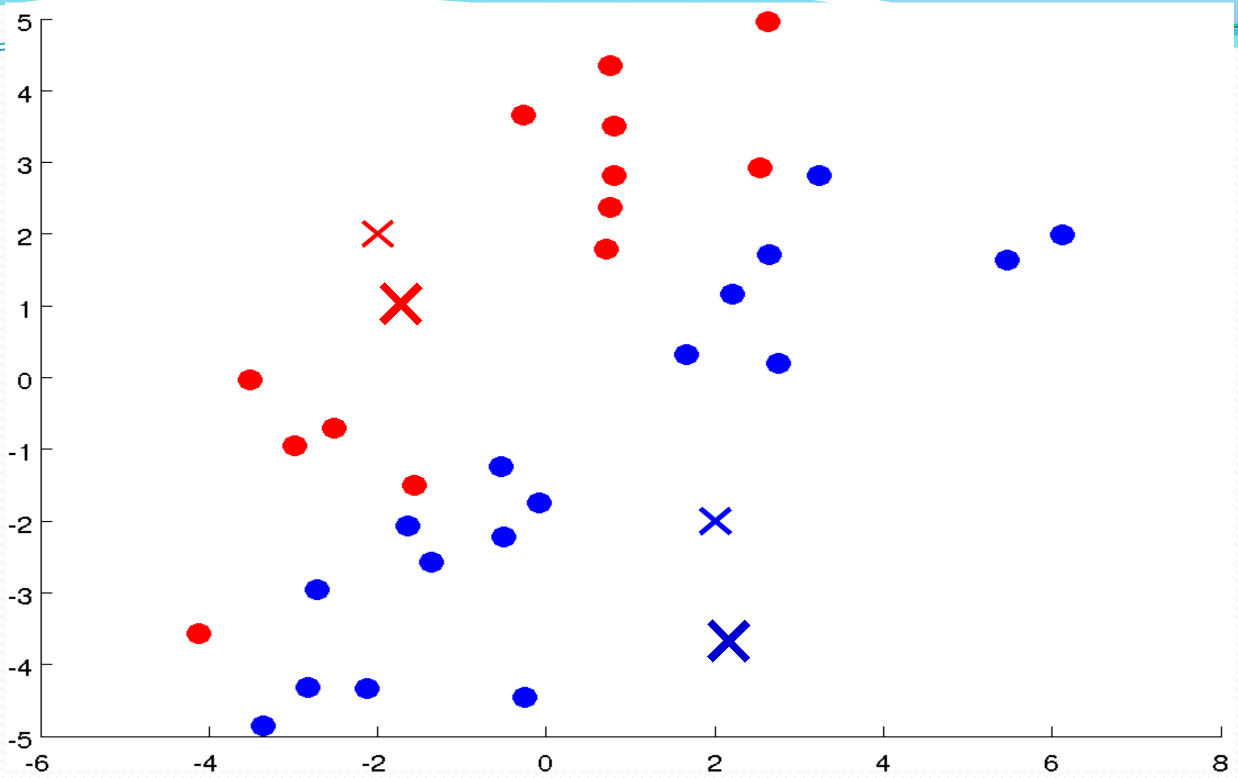

➢In the cluster assignment step, the algorithm goes through each of the training examples, and depending on whether it's closer to which cluster centroid, it is going to assign each of the data points to one of the two cluster centroids.

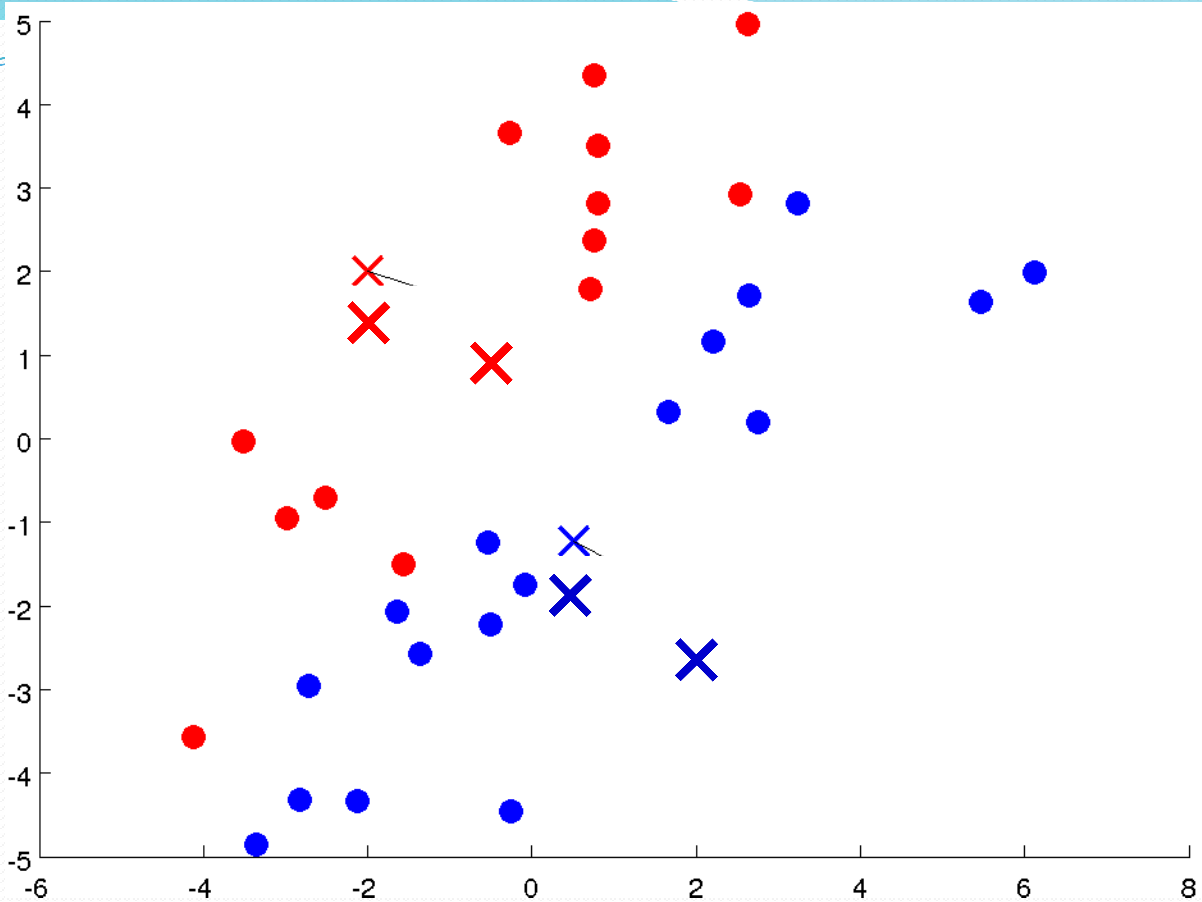

➢In the move centroid step, the algorithm is going to move the two cluster centroids to the average of the points of the same cluster.

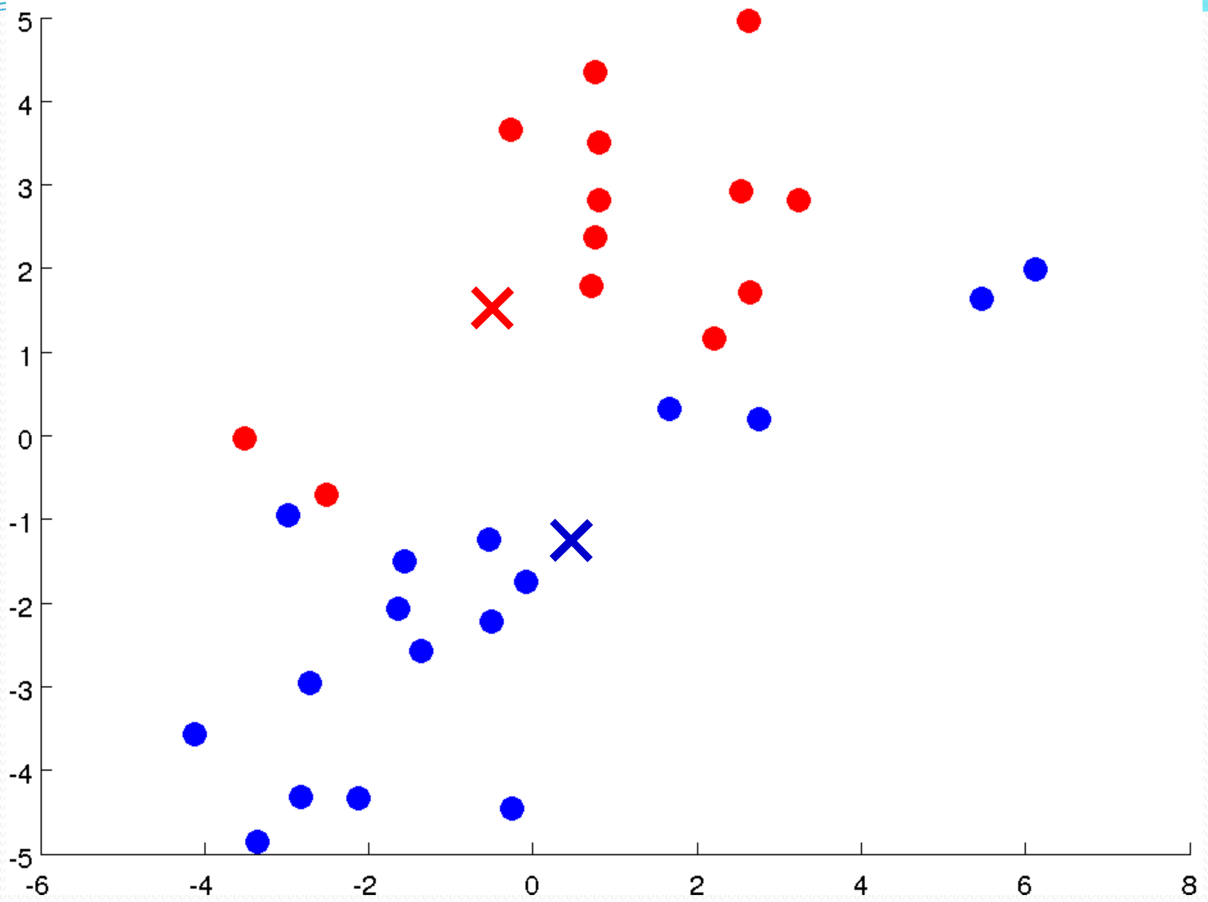
5
4
3
2
1
0
-1
-2
-3
-4
-5
-6
-4
-2
0
2
4
6
8

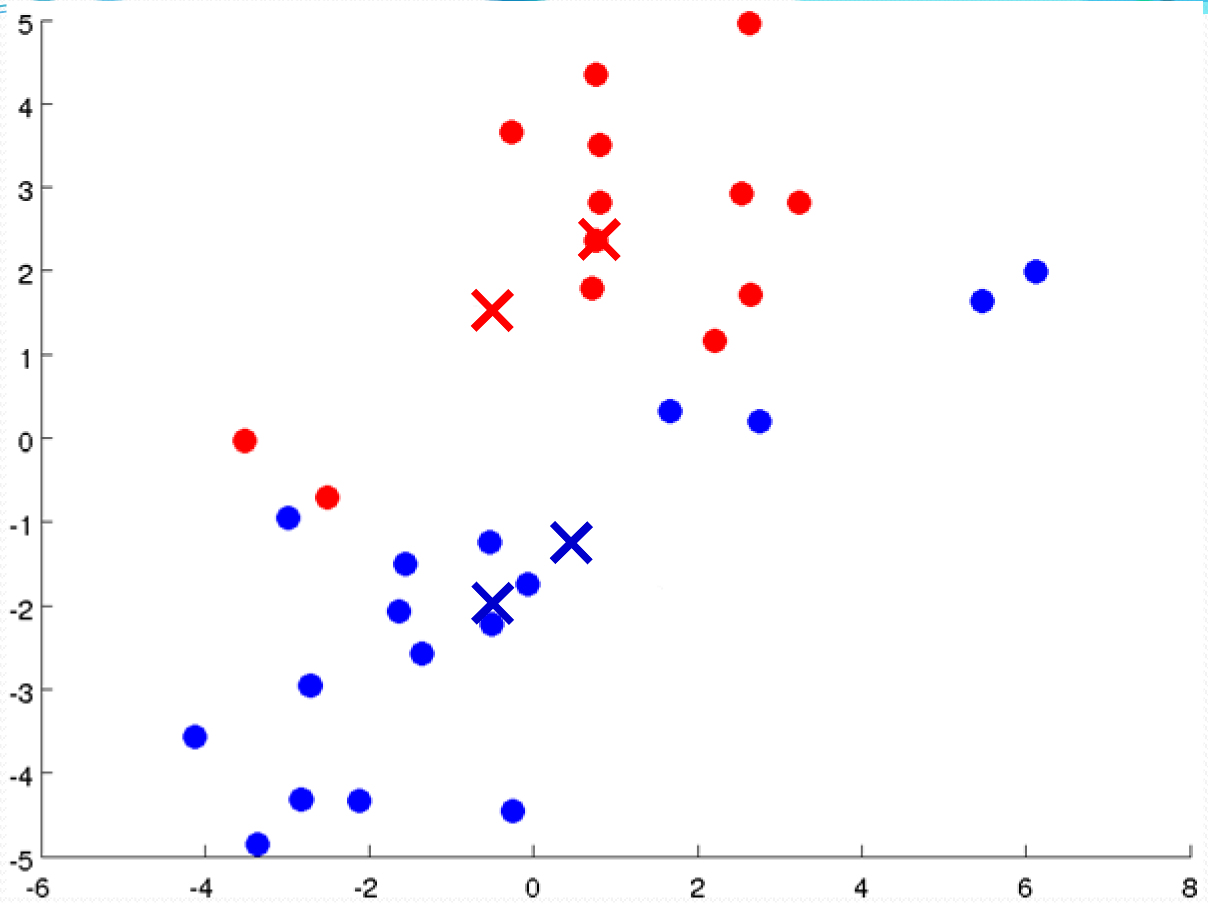
5
4
3
2
1
0
-1
-2
-3
-4
-5
-6
-4
-2
0
2
4
6
8

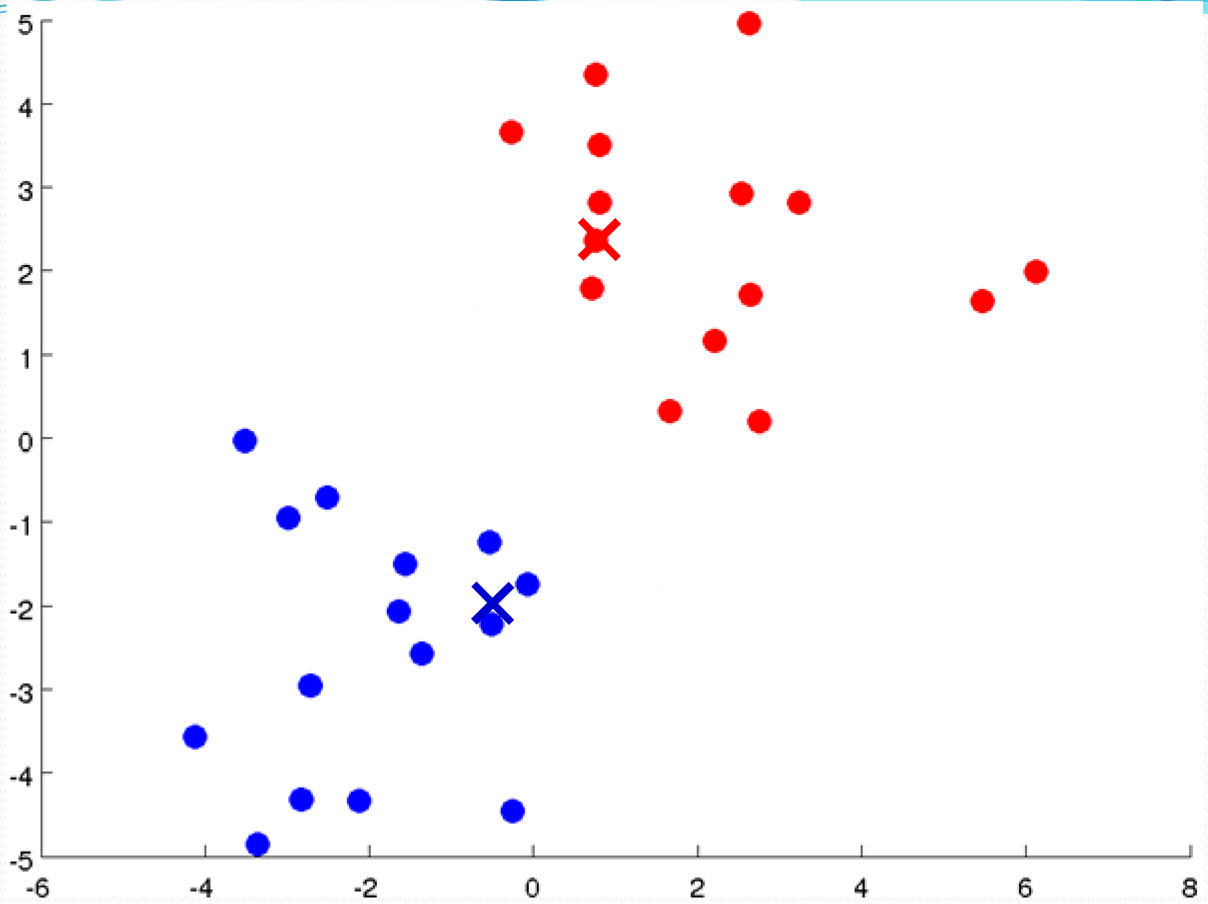
5
4
3
2
1
0
-1
-2
-3
-4
-5
-6
-4
-2
0
2
4
6
8

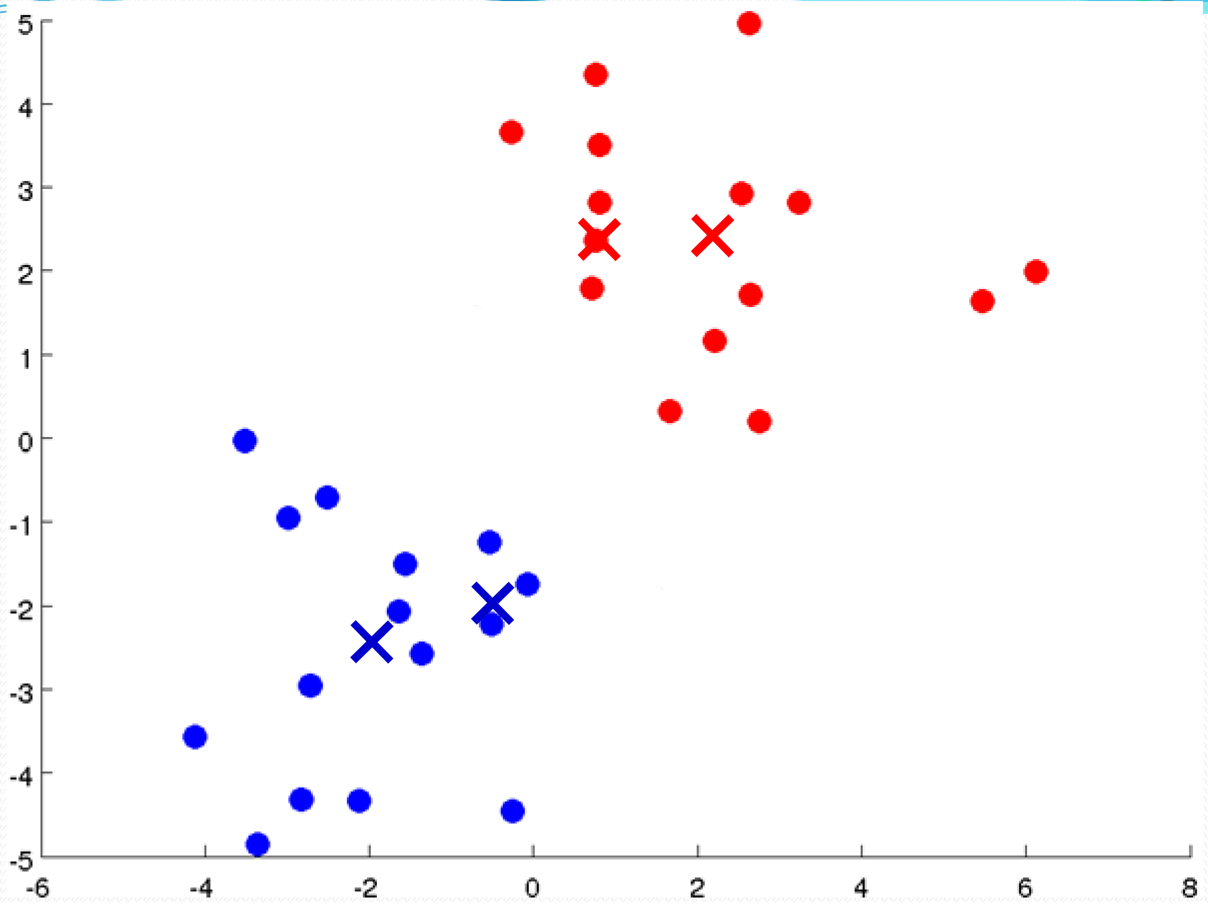
5
4
3
2
1
0
-1
-2
-3
-4
-5
-6
-4
-2
0
2
4
6
8

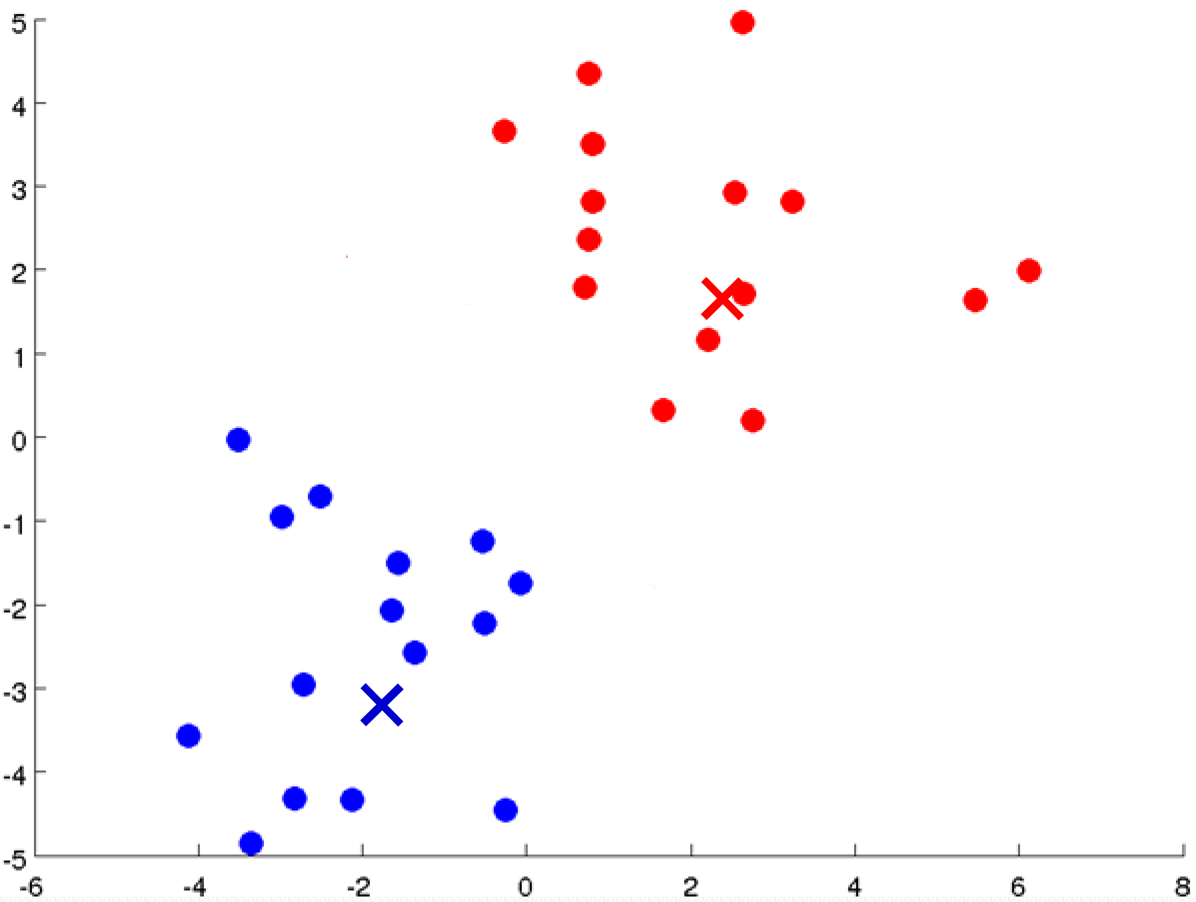

➢If you keep running additional iterations of K means from here the cluster centroids will not change any further, so at this point, K means has converged

K-means algorithm

Input:

- K (number of clusters)
- Training set $\{x^{(1)}, x^{(2)}, \ldots, x^{(m)}\}$

$x^{(i)} \in \mathbb{R}^n$ (drop $x_0 = 1$ convention)

K-means algorithm

Randomly initialize K cluster centroids $\mu_1, \mu_2, \ldots, \mu_K \in \mathbb{R}^n$

Repeat {

for i = 1 to m

$c^{(i)}$:= index (from 1 to K) of cluster centroid closest to $x^{(i)}$ $\min_k \; ||x^{(i)} - \mu_k||^2$

for k = 1 to K

μ_k := average (mean) of points assigned to cluster k

}

K-means for non-separated clusters

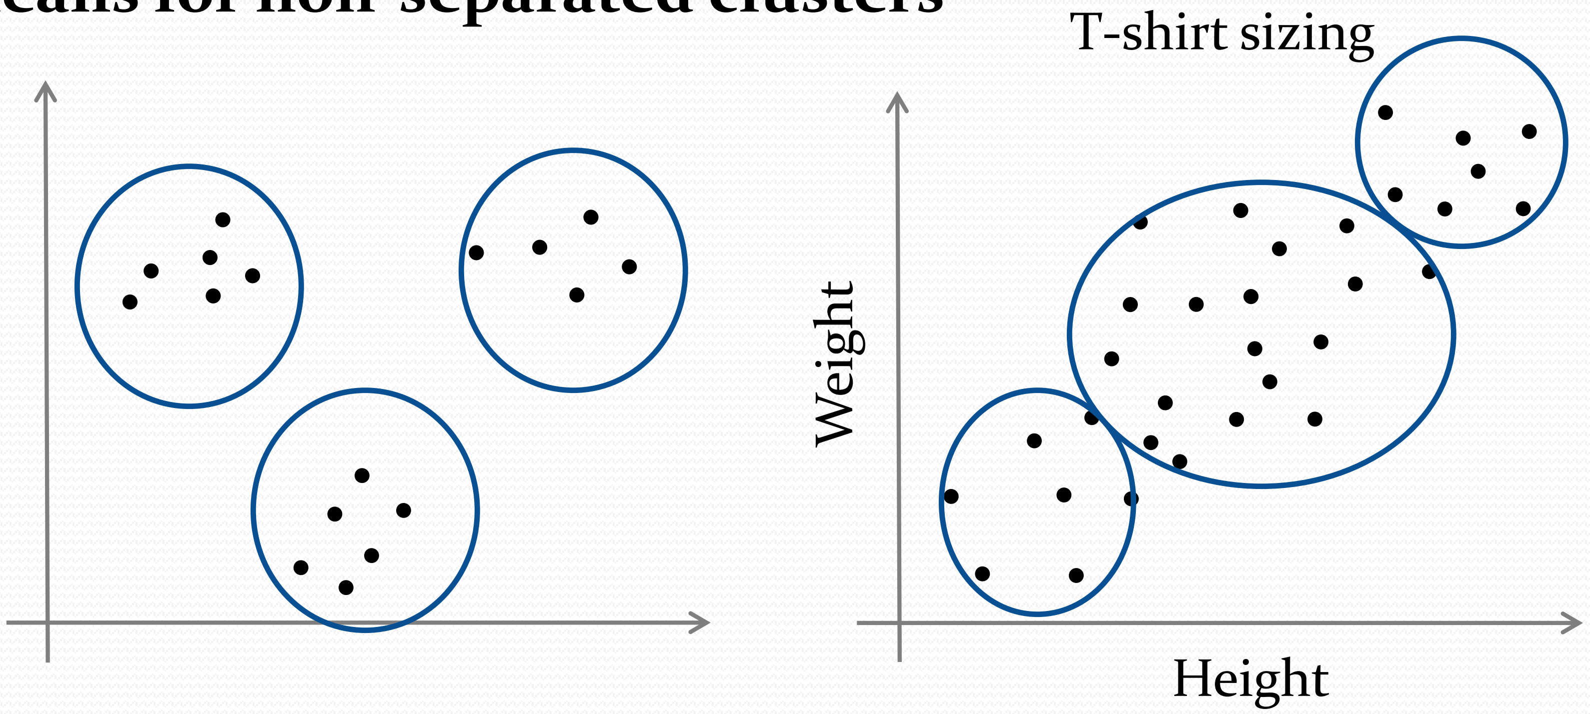

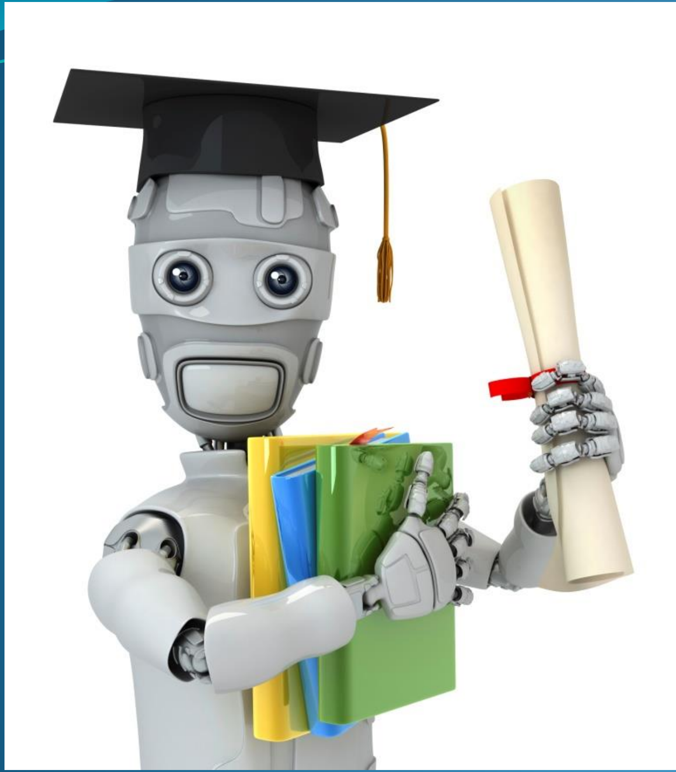

Machine Learning

Clustering

Optimization objective

K-means optimization objective

$c^{(i)}$ = index of cluster (1,2,…,K) to which example $x^{(i)}$ is currently assigned

μ_k = cluster centroid k ($\mu_k \in \mathbb{R}^n$)

$\mu_{c^{(i)}}$ = cluster centroid of cluster to which example $x^{(i)}$ has been assigned

Optimization objective:

$$J(c^{(1)}, \dots, c^{(m)}, \mu_1, \dots, \mu_K) = \frac{1}{m} \sum_{i=1}^{m} ||x^{(i)} - \mu_{c^{(i)}}||^2$$

$$\min_{\substack{c^{(1)}, \dots, c^{(m)}, \\ \mu_1, \dots, \mu_K}} J(c^{(1)}, \dots, c^{(m)}, \mu_1, \dots, \mu_K)$$

K-means algorithm

Randomly initialize K cluster centroids $\mu_1, \mu_2, \ldots, \mu_K \in \mathbb{R}^n$

Repeat {

for i = 1 to m

$c^{(i)}$:= index (from 1 to K) of cluster centroid closest to $x^{(i)}$

for k = 1 to K

μ_k := average (mean) of points assigned to k cluster

}

Random initialization

- Should have $K < m$
- Randomly pick K training examples.
- Set $\mu_1, \ldots, \mu_K$ equal to these K examples.
- By these two illustrations on the right. You might really guess that K-means can end up converging to different solutions depending on exactly how the clusters were initialized.

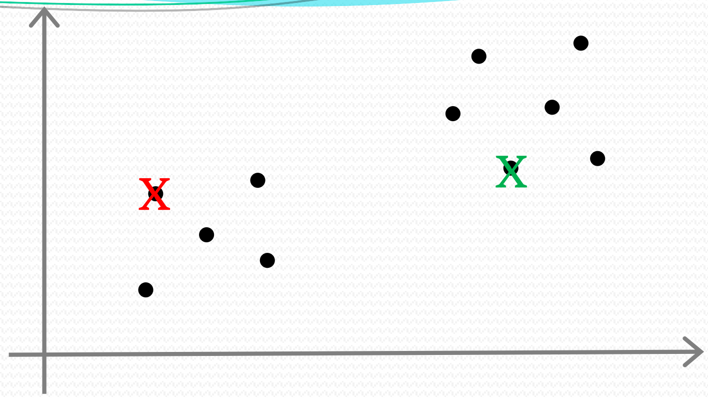

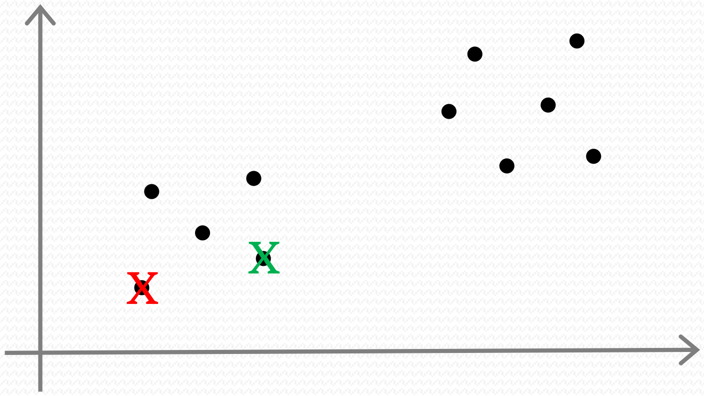

Local optima

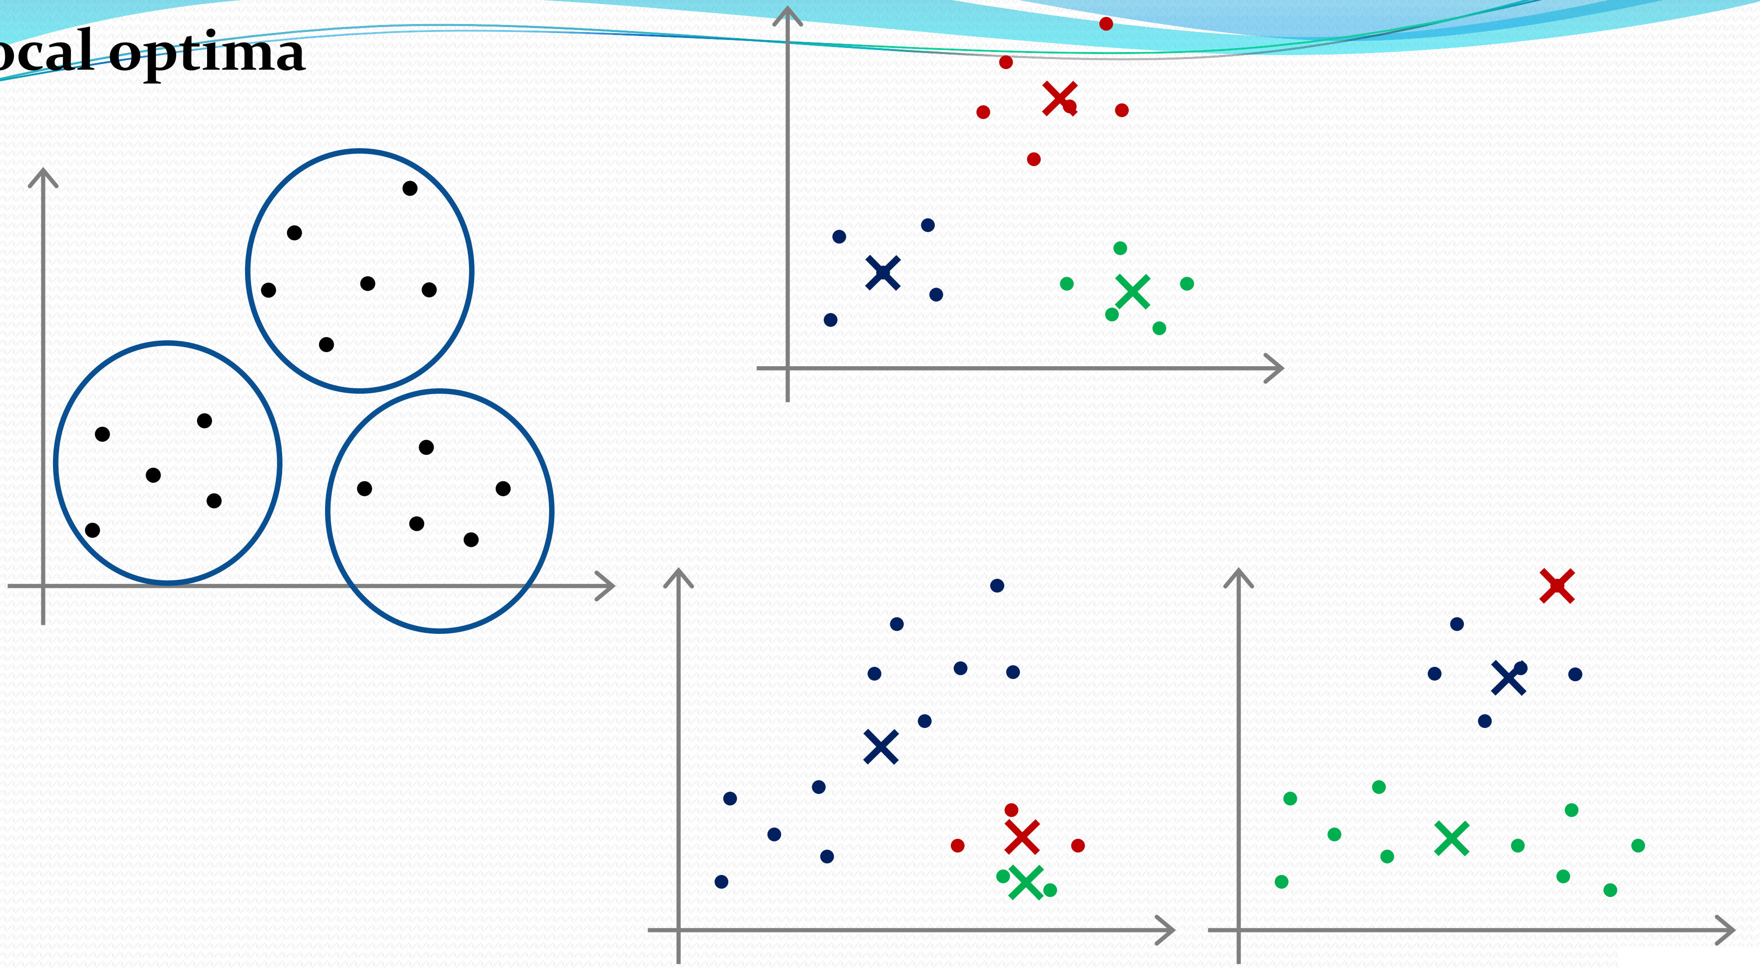

Random initialization

For i = 1 to 100 {

Randomly initialize K-means.

Run K-means. Get $c^{(1)}, \ldots, c^{(m)}, \mu_1, \ldots, \mu_K$.

Compute cost function

$$J(c^{(1)}, \ldots, c^{(m)}, \mu_1, \ldots, \mu_K)$$

}

Pick clustering that gave lowest cost $J(c^{(1)}, \ldots, c^{(m)}, \mu_1, \ldots, \mu_K)$

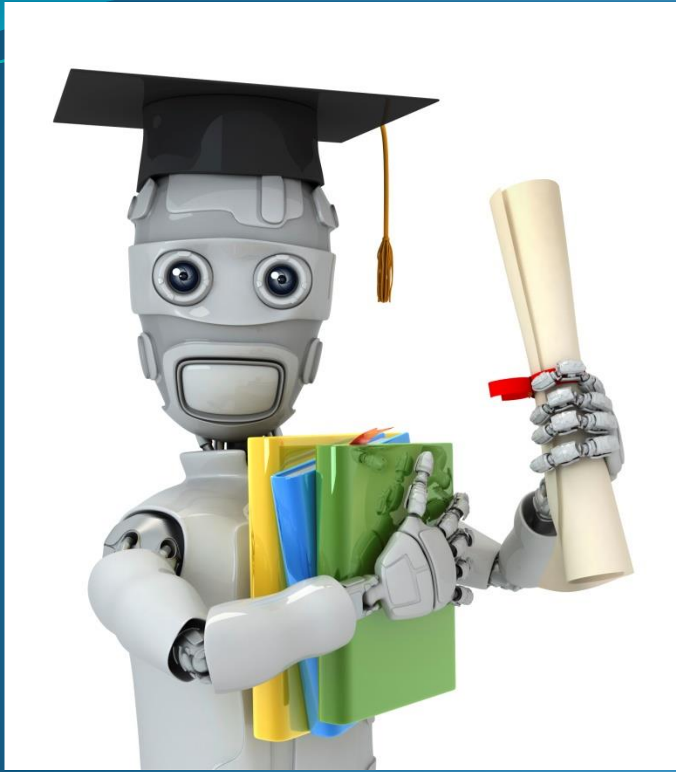

Machine Learning

Clustering

Choosing the number of clusters

What is the right value of K?

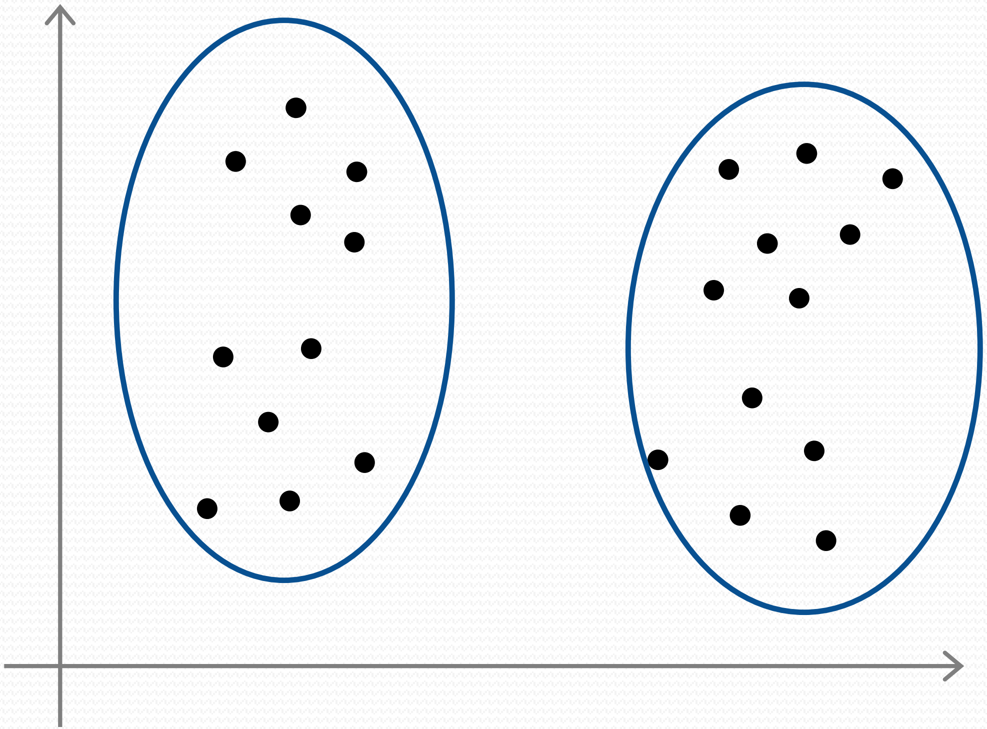

Choosing the value of K

Elbow method:

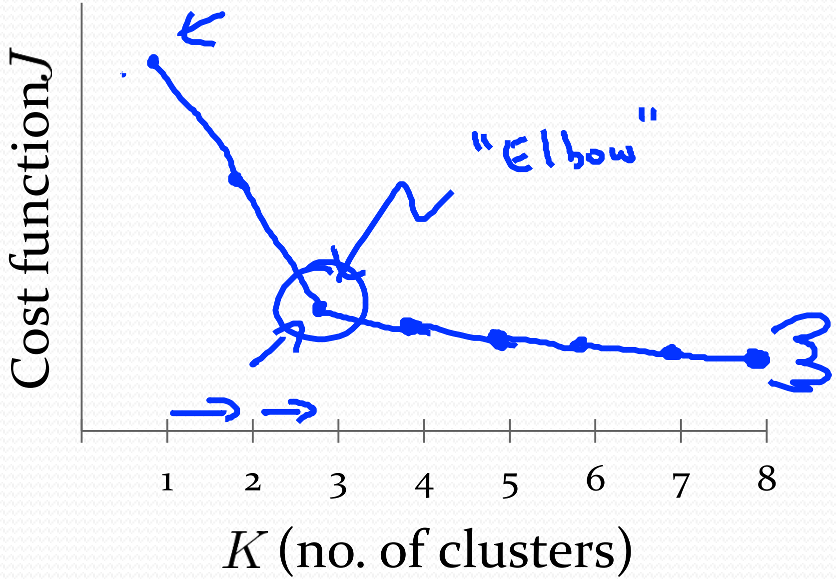

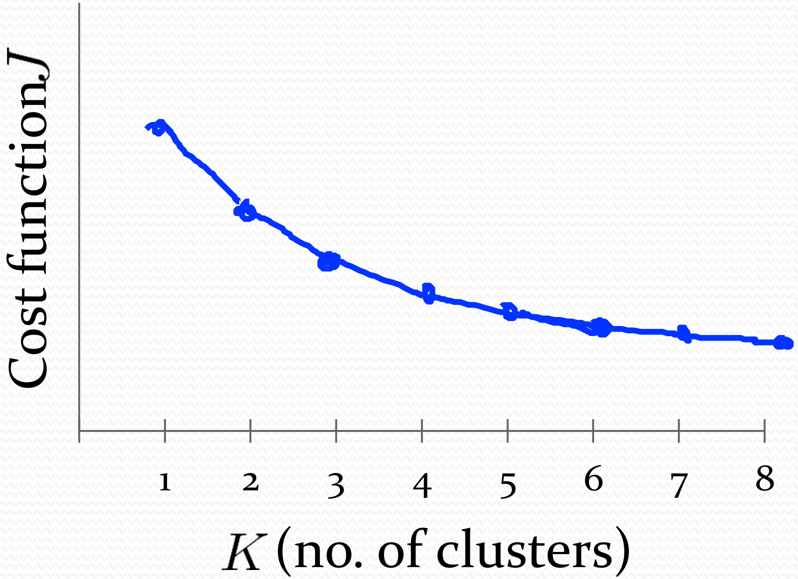

Choosing the value of K

Sometimes, you're running K-means to get clusters to use for some later/downstream purpose. Evaluate K-means based on a metric for how well it performs for that later purpose.

E.g.

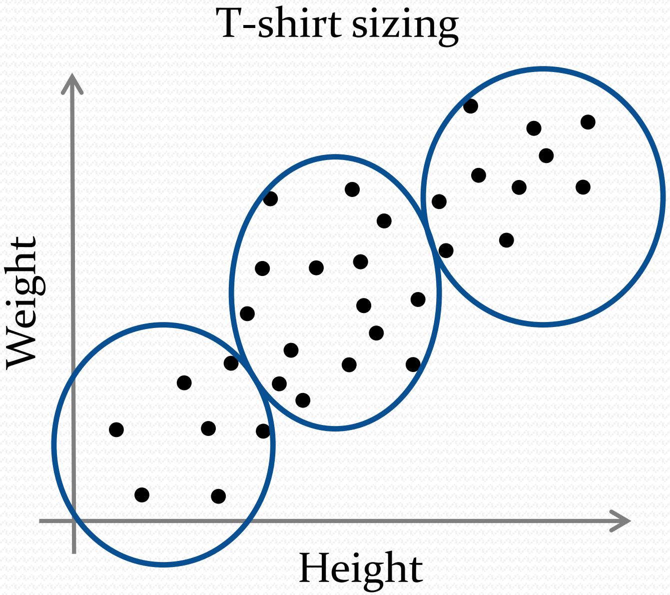

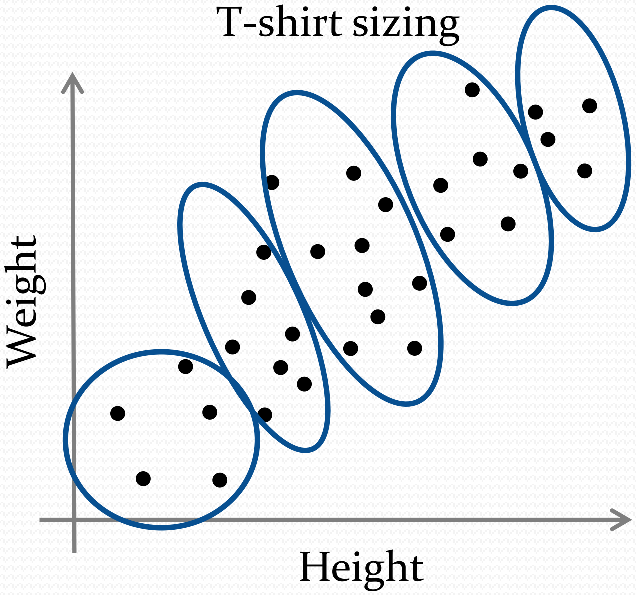

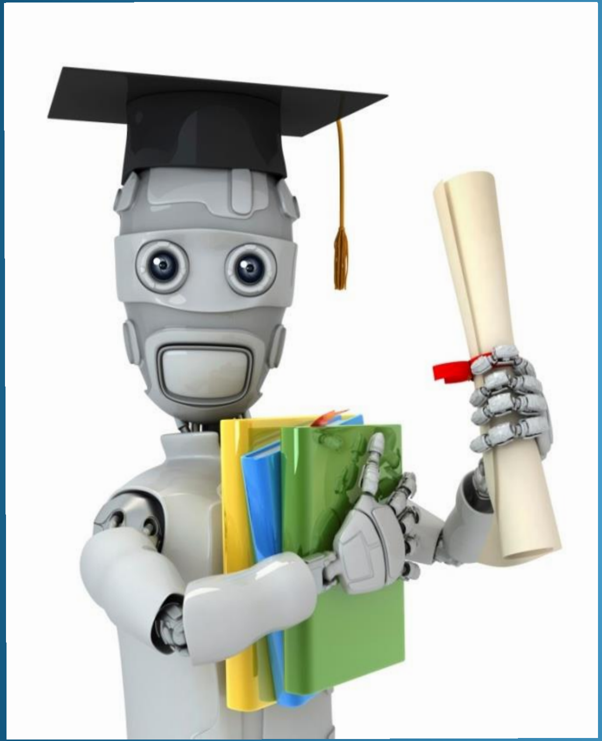

Thanks